arXiv:1703.00441v1 [cs.LG] 1 Mar 2017

Learning to Optimize Neural Nets

Ke Li [1] **Jitendra Malik** [1]

Abstract

Learning to Optimize (Li & Malik, 2016) is a recently proposed framework for learning optimization algorithms using reinforcement learning. In this paper, we explore learning an optimization algorithm for training shallow neural nets. Such high-dimensional stochastic optimization problems present interesting challenges for existing reinforcement learning algorithms. We develop an extension that is suited to learning optimization algorithms in this setting and demonstrate that the learned optimization algorithm consistently outperforms other known optimization algorithms even on unseen tasks and is robust to changes in stochasticity of gradients and the neural net architecture. More specifically, we show that an optimization algorithm trained with the proposed method on the problem of training a neural net on MNIST generalizes to the problems of training neural nets on the Toronto Faces Dataset, CIFAR-10 and CIFAR-100.

1. Introduction

Machine learning is centred on the philosophy that learning patterns automatically from data is generally better than meticulously crafting rules by hand. This data-driven approach has delivered: today, machine learning techniques can be found in a wide range of application areas, both in AI and beyond. Yet, there is one domain that has conspicuously been left untouched by machine learning: the design of tools that power machine learning itself.

One of the most widely used tools in machine learning is optimization algorithms. We have grown accustomed to seeing an optimization algorithm as a black box that takes in a model that we design and the data that we collect and outputs the optimal model parameters. The optimization algorithm itself largely stays static: its design is reserved for human experts, who must toil through many rounds of theoretical analysis and empirical validation to devise a better optimization algorithm. Given this state of affairs, perhaps it is time for us to start practicing what we preach and learn how to learn.

Recently, Li & Malik (2016) and Andrychowicz et al. (2016) introduced two different frameworks for learning optimization algorithms. Whereas Andrychowicz et al. (2016) focuses on learning an optimization algorithm for training models on a particular task, Li & Malik (2016) sets a more ambitious objective of learning an optimization algorithm for training models that is task-independent. We study the latter paradigm in this paper and develop a method for learning an optimization algorithm for high-dimensional stochastic optimization problems, like the problem of training shallow neural nets.

Under the "Learning to Optimize" framework proposed by Li & Malik (2016), the problem of learning an optimization algorithm is formulated as a reinforcement learning problem. We consider the general structure of an unconstrained continuous optimization algorithm, as shown in Algorithm 1. In each iteration, the algorithm takes a step Δx and uses it to update the current iterate $x^{(i)}$. In hand-engineered optimization algorithms, Δx is computed using some fixed formula ϕ that depends on the objective function, the current iterate and past iterates. Often, it is simply a function of the current and past gradients.

Algorithm 1 General structure of optimization algorithms

Require: Objective function f
  $x^{(0)} \leftarrow$ random point in the domain of f
  for $i = 1, 2, \ldots$ **do**
    $\Delta x \leftarrow \phi(f, \{x^{(0)}, \ldots, x^{(i-1)}\})$
    if stopping condition is met **then**
      return $x^{(i-1)}$
    end if
    $x^{(i)} \leftarrow x^{(i-1)} + \Delta x$
  end for

Different choices of ϕ yield different optimization algorithms and so each optimization algorithm is essentially characterized by its update formula ϕ. Hence, by learning ϕ, we can learn an optimization algorithm. Li & Malik (2016) observed that an optimization algorithm can be viewed as a Markov decision process (MDP), where the state includes the current iterate, the action is the step vec-

[1]University of California, Berkeley, CA 94720, United States. Correspondence to: Ke Li <ke.li@eecs.berkeley.edu>.

tor Δx and the policy is the update formula ϕ. Hence, the problem of learning ϕ simply reduces to a policy search problem.

In this paper, we build on the method proposed in (Li & Malik, 2016) and develop an extension that is suited to learning optimization algorithms for high-dimensional stochastic problems. We use it to learn an optimization algorithm for training shallow neural nets and show that it outperforms popular hand-engineered optimization algorithms like ADAM (Kingma & Ba, 2014), AdaGrad (Duchi et al., 2011) and RMSprop (Tieleman & Hinton, 2012) and an optimization algorithm learned using the supervised learning method proposed in (Andrychowicz et al., 2016). Furthermore, we demonstrate that our optimization algorithm learned from the experience of training on MNIST generalizes to training on other datasets that have very dissimilar statistics, like the Toronto Faces Dataset, CIFAR-10 and CIFAR-100.

2. Related Work

The line of work on learning optimization algorithms is fairly recent. Li & Malik (2016) and Andrychowicz et al. (2016) were the first to propose learning optimization algorithms. Li & Malik (2016) explored learning task-independent optimization algorithms for training various models and used reinforcement learning to learn the optimization algorithm, while Andrychowicz et al. (2016) investigated learning task-dependent optimization algorithms and used supervised learning.

More broadly related is work on "meta-learning" and "learning to learn" (Baxter et al., 1995; Vilalta & Drissi, 2002; Brazdil et al., 2008; Thrun & Pratt, 2012). Historically, these terms have been used by different authors to refer to approaches that aim to achieve different objectives, which essentially differ in the type of knowledge learned at the meta-level. Many methods try to learn commonalities shared across a family of related tasks; this line of work has blossomed into an area that has later become known as transfer learning and multi-task learning. Other methods (Brazdil et al., 2003; Schmidhuber, 2004) aim to learn how to choose a base-level learner that would perform the best on a given task. One challenge under this setting is to decide on a parameterization of the space of base-level learners that is both rich enough to be capable of representing disparate base-level learners and compact enough to permit tractable search over this space. Brazdil et al. (2003) proposes a nonparametric representation and stores examples of different base-level learners in a database, whereas Schmidhuber (2004) proposes representing base-level learners as general-purpose programs. The former has limited representation power, while the latter makes search and learning in the space of base-level learners intractable.

Work on learning optimization algorithms can be viewed as a different approach to learning to learn/meta-learning, where the process of training a base-level learner is represented as a sequence of updates to the learner's parameters.

Because the proposed method learns an algorithm, it is related to the line of work on program induction, the goal of which is to learn programs from examples of input and output. Various approaches have been explored: Cramer (1985) investigated evolving programs represented as abstract syntax trees using genetic algorithms, Liang et al. (2010) defines a hierarchical Bayesian prior over programs in a formal language and infers the best program using approximate sampling and neural Turing machines (Graves et al., 2014) try to predict output examples from input examples directly. Hochreiter et al. (2001) observed that a recurrent neural net with a particular setting of weights that takes a training example as input at each time step can be viewed as an online learning algorithm and learns an online learning algorithm by learning the weights of the recurrent neural net.

Work on hyperparameter optimization tries to design automatic ways of searching for the best hyperparameters for training a model on a particular task. Various kinds of methods have been proposed, such those based on Bayesian optimization (Hutter et al., 2011; Bergstra et al., 2011; Snoek et al., 2012; Swersky et al., 2013; Feurer et al., 2015), random search (Bergstra & Bengio, 2012) and gradient-based optimization (Bengio, 2000; Domke, 2012; Maclaurin et al., 2015). Note that the discovered hyperparameters are generally specific to the task and the model, and hyperparameter optimization must be rerun for a new task or model.

More closely related are methods for online hyperparameter adaptation, which adjust particular types of hyperparameters automatically while performing optimization. Some rules for adjusting hyperparameters are manually designed (Bray et al., 2004), while others learn policies for adjusting the step size (Hansen, 2016; Daniel et al., 2016; Fu et al., 2016) or the damping factor in the Levenberg-Marquardt algorithm (Ruvolo et al., 2009). Unlike this line of work, this paper explores learning a general optimization algorithm, which may differ from existing optimization algorithms both in direction and size of the step taken in each iteration. A different line of work (Gregor & LeCun, 2010; Sprechmann et al., 2013) parameterizes intermediate operands of special-purpose solvers for a class of optimization problems that arise in sparse coding and learns them using supervised learning.

3. Learning to Optimize

3.1. Setting

In the "Learning to Optimize" framework, we are given a set of training objective functions $f_1, \ldots, f_n$ drawn from some distribution $\mathcal{F}$. An optimization algorithm $\mathcal{A}$ takes an objective function f and an initial iterate $x^{(0)}$ as input and produces a sequence of iterates $x^{(1)}, \ldots, x^{(T)}$, where $x^{(T)}$ is the solution found by the optimizer. We are also given a distribution $\mathcal{D}$ that generates the initial iterate $x^{(0)}$ and a meta-loss $\mathcal{L}$, which takes an objective function f and a sequence of iterates $x^{(1)}, \ldots, x^{(T)}$ produced by an optimization algorithm as input and outputs a scalar that measures the quality of the iterates. The goal is to learn an optimization algorithm $\mathcal{A}^*$ such that $\mathbb{E}_{f\sim\mathcal{F}, x^{(0)}\sim\mathcal{D}}\left[\mathcal{L}(f, \mathcal{A}^*(f, x^{(0)}))\right]$ is minimized. The meta-loss is chosen to penalize optimization algorithms that exhibit behaviours we find undesirable, like slow convergence or excessive oscillations. Assuming we would like to learn an algorithm that minimizes the objective function it is given, a good choice of meta-loss would then simply be $\sum_{i=1}^{T} f(x^{(i)})$, which can be interpreted as the area under the curve of objective values over time.

The objective functions $f_1, \ldots, f_n$ may correspond to loss functions for training base-level learners, in which case the algorithm that learns the optimization algorithm can be viewed as a meta-learner. In this setting, each objective function is the loss function for training a particular base-learner on a particular task, and so the set of training objective functions can be loss functions for training a base-learner or a family of base-learners on different tasks. At test time, the learned optimization algorithm is evaluated on unseen objective functions, which correspond to loss functions for training base-learners on new tasks, which may be completely unrelated to tasks used for training the optimization algorithm. Therefore, the learned optimization algorithm must not learn anything about the tasks used for training. Instead, the goal is to learn an optimization algorithm that can exploit the geometric structure of the error surface induced by the base-learners. For example, if the base-level model is a neural net with ReLU activation units, the optimization algorithm should hopefully learn to leverage the piecewise linearity of the model. Hence, there is a clear division of responsibilities between the meta-learner and base-learners. The knowledge learned at the meta-level should be pertinent for all tasks, whereas the knowledge learned at the base-level should be task-specific. The meta-learner should therefore generalize across tasks, whereas the base-learner should generalize across instances.

3.2. RL Preliminaries and Formulation

The goal of reinforcement learning is to learn to interact with an environment in a way that minimizes cumulative costs that are expected to be incurred over time. The environment is formalized as a partially observable Markov decision process (POMDP)[1], which is defined by the tuple $(\mathcal{S}, \mathcal{O}, \mathcal{A}, p_i, p, p_o, c, T)$, where $\mathcal{S} \subseteq \mathbb{R}^D$ is the set of states, $\mathcal{O} \subseteq \mathbb{R}^{D'}$ is the set of observations, $\mathcal{A} \subseteq \mathbb{R}^d$ is the set of actions, $p_i\left(s_0\right)$ is the probability density over initial states s_0, $p\left(s_{t+1} \,|s_t, a_t\right)$ is the probability density over the subsequent state s_{t+1} given the current state s_t and action a_t, $p_o\left(o_t \,|s_t\right)$ is the probability density over the current observation o_t given the current state s_t, $c : \mathcal{S} \to \mathbb{R}$ is a function that assigns a cost to each state and T is the time horizon. Often, the probability densities p_i, p and p_o are not given explicitly, but can be accessed indirectly via sampling.

In our setting, the state s_t consists of the current iterate $x^{(t)}$ and features $\Phi(\cdot)$ that depend on the history of iterates $x^{(1)}, \ldots, x^{(t)}$, (possibly noisy) gradients $\nabla\hat{f}(x^{(1)}), \ldots, \nabla\hat{f}(x^{(t)})$ and (possibly noisy) objective values $\hat{f}(x^{(1)}), \ldots, \hat{f}(x^{(t)})$. The observation o_t excludes $x^{(t)}$ and consists of possibly different features $\Psi\left(\left\{x^{(i)}, \nabla\hat{f}(x^{(i)}), \hat{f}(x^{(i)})\right\}_{i=0}^{t}\right)$. The action a_t is the step Δx that will be used to update the iterate. The initial probability density p_i is defined implicitly in terms of $\mathcal{D}$ and is the density of the random variable $\left(x^{(0)}, \Phi\left(x^{(0)}, \nabla\hat{f}(x^{(0)}), \hat{f}(x^{(0)})\right)\right)^T$ where $x^{(0)} \sim \mathcal{D}$ and $f \sim \mathcal{F}$. The transition probability density p is the density of the random variable $\left(x^{(t)} + \Delta x, \Phi\left(\left\{x^{(i)}, \nabla\hat{f}(x^{(i)}), \hat{f}(x^{(i)})\right\}_{i=0}^{t+1}\right)\right)^T$ given $x^{(t)}$, $\Phi\left(\left\{x^{(i)}, \nabla\hat{f}(x^{(i)}), \hat{f}(x^{(i)})\right\}_{i=0}^{t}\right)$ and Δx, where $f \sim \mathcal{F}$. The observation probability density p_o is also defined implicitly and depends on the definition of Φ and Ψ. Assuming the goal is to learn an optimization algorithm that minimizes the objective function, the cost c of a state $s_t = \left(x^{(t)}, \Phi\left(\cdot\right)\right)^T$ is simply the true objective value $f(x^{(t)})$.

A policy $\pi\left(a_t \,|o_t, t\right)$ is a conditional probability density over actions a_t given the current observation o_t and time step t. When a policy is independent of t, it is known as a stationary policy. The goal of the reinforcement learning algorithm is to learn a policy π^* that minimizes the total expected cost over time. More precisely,

$$\pi^* = \arg\min_{\pi} \mathbb{E}_{s_0, a_0, s_1, \ldots, s_T}\left[\sum_{t=0}^{T} c(s_t)\right],$$

where the expectation is taken with respect to the joint distribution over the sequence of states and actions, often re-

[1] What is described is an undiscounted finite-horizon POMDP with continuous state, observation and action spaces.

ferred to as a trajectory, which has the density

$$q(s_0, a_0, s_1, \ldots, s_T) = \int_{o_0,\ldots,o_T} p_i(s_0)\, p_o(o_0|\, s_0) \prod_{t=0}^{T-1} \pi(a_t|\, o_t, t)\, p(s_{t+1}|\, s_t, a_t)\, p_o(o_{t+1}|\, s_{t+1}).$$

To make learning tractable, π is often constrained to lie in a parameterized family. A common assumption is that $\pi(a_t|\, o_t, t) = \mathcal{N}(\mu^\pi(o_t), \Sigma^\pi(o_t))$, where $\mathcal{N}(\mu, \Sigma)$ denotes the density of a Gaussian with mean μ and covariance Σ. The functions $\mu^\pi(\cdot)$ and possibly $\Sigma^\pi(\cdot)$ are modelled using function approximators, whose parameters are learned.

In our setting, any particular policy $\pi(a_t\,|o_t, t)$, which generates $a_t = \Delta x$ given $o_t = \Phi(\cdot)$ at every time step corresponds to a particular (noisy) update formula ϕ, and therefore a particular (noisy) optimization algorithm. In practice, stochasticity is only used for training the policy; when testing the policy, $\Sigma^\pi(o_t)$ is typically set to zero, thereby making the optimization algorithm behave deterministically. Therefore, learning an optimization algorithm simply reduces to searching for the optimal policy.

3.3. Guided Policy Search

The reinforcement learning method we use is guided policy search (GPS) (Levine et al., 2015), which is a policy search method designed for searching over large classes of expressive non-linear policies in continuous state and action spaces. It maintains two policies, ψ and π, where the former lies in a time-varying linear policy class in which the optimal policy can found in closed form, and the latter lies in a stationary non-linear policy class in which policy optimization is challenging. In each iteration, it performs policy optimization on ψ, and uses the resulting policy as supervision to train π.

More precisely, GPS solves the following constrained optimization problem:

$$\min_{\theta,\eta} \mathbb{E}_\psi\left[\sum_{t=0}^{T} c(s_t)\right] \text{ s.t. } \psi(a_t|\, s_t, t; \eta) = \pi(a_t|\, s_t; \theta)\ \forall a_t, s_t, t$$

where η and θ denote the parameters of ψ and π respectively, $\mathbb{E}_\rho[\cdot]$ denotes the expectation taken with respect to the trajectory induced by a policy ρ and $\pi(a_t|\, s_t; \theta) := \int_{o_t} \pi(a_t|\, o_t; \theta)\, p_o(o_t|\, s_t)$[2].

Since there are an infinite number of equality constraints, the problem is relaxed by enforcing equality on the mean actions taken by ψ and π at every time step[3]. So, the problem becomes:

$$\min_{\theta,\eta} \mathbb{E}_\psi\left[\sum_{t=0}^{T} c(s_t)\right] \text{ s.t. } \mathbb{E}_\psi[a_t] = \mathbb{E}_\psi[\mathbb{E}_\pi[a_t|\, s_t]]\ \forall t$$

This problem is solved using Bregman ADMM (Wang & Banerjee, 2014), which performs the following updates in each iteration:

$$\eta \leftarrow \arg\min_\eta \sum_{t=0}^{T} \mathbb{E}_\psi\left[c(s_t) - \lambda_t^T a_t\right] + \nu_t D_t(\eta, \theta)$$

$$\theta \leftarrow \arg\min_\theta \sum_{t=0}^{T} \lambda_t^T \mathbb{E}_\psi[\mathbb{E}_\pi[a_t|\, s_t]] + \nu_t D_t(\theta, \eta)$$

$$\lambda_t \leftarrow \lambda_t + \alpha\nu_t(\mathbb{E}_\psi[\mathbb{E}_\pi[a_t|\, s_t]] - \mathbb{E}_\psi[a_t])\ \forall t,$$

where $D_t(\theta, \eta) := \mathbb{E}_\psi[D_{KL}(\pi(a_t|\, s_t; \theta) \| \psi(a_t|\, s_t, t; \eta))]$ and $D_t(\eta, \theta) := \mathbb{E}_\psi[D_{KL}(\psi(a_t|\, s_t, t; \eta) \| \pi(a_t|\, s_t; \theta))]$.

The algorithm assumes that $\psi(a_t|\, s_t, t; \eta) = \mathcal{N}(K_t s_t + k_t, G_t)$, where $\eta := (K_t, k_t, G_t)_{t=1}^T$ and $\pi(a_t|\, o_t; \theta) = \mathcal{N}(\mu_\omega^\pi(o_t), \Sigma^\pi)$, where $\theta := (\omega, \Sigma^\pi)$ and $\mu_\omega^\pi(\cdot)$ can be an arbitrary function that is typically modelled using a nonlinear function approximator like a neural net.

At the start of each iteration, the algorithm constructs a model of the transition probability density $\tilde{p}(s_{t+1}|\, s_t, a_t, t; \zeta) = \mathcal{N}(A_t s_t + B_t a_t + c_t, F_t)$, where $\zeta := (A_t, B_t, c_t, F_t)_{t=1}^T$ is fitted to samples of s_t drawn from the trajectory induced by ψ, which essentially amounts to a local linearization of the true transition probability $p(s_{t+1}|\, s_t, a_t, t)$. We will use $\mathbb{E}_{\tilde{\psi}}[\cdot]$ to denote expectation taken with respect to the trajectory induced by ψ under the modelled transition probability $\tilde{p}$. Additionally, the algorithm fits local quadratic approximations to $c(s_t)$ around samples of s_t drawn from the trajectory induced by ψ so that $c(s_t) \approx \tilde{c}(s_t) := \frac{1}{2} s_t^T C_t s_t + d_t^T s_t + h_t$ for s_t's that are near the samples.

With these assumptions, the subproblem that needs to be solved to update $\eta = (K_t, k_t, G_t)_{t=1}^T$ becomes:

$$\min_\eta \mathbb{E}_{\tilde{\psi}}\left[\sum_{t=0}^{T} \tilde{c}(s_t) - a_t^T \lambda_t\right] + \nu_t D(\eta, \theta)$$

$$\text{s.t. } \mathbb{E}_{\tilde{\psi}}\left[\sum_{t=0}^{T} D_{KL}(\psi(a_t|\, s_t, t; \eta) \| \psi(a_t|\, s_t, t; \eta'))\right] \le \epsilon,$$

where η' denotes the old η from the previous iteration. Because $\tilde{p}$ and $\tilde{c}$ are only valid locally around the trajectory induced by ψ, the constraint is added to limit the amount by which η is updated. It turns out that the unconstrained problem can be solved in closed form using a dynamic programming algorithm known as linear-quadratic-Gaussian (LQG)

[2] In practice, the explicit form of the observation probability p_o is usually not known or the integral may be intractable to compute. So, a linear Gaussian model is fitted to samples of s_t and a_t and used in place of the true $\pi(a_t|\, s_t; \theta)$ where necessary.

[3] Though the Bregman divergence penalty is applied to the original probability distributions over a_t.

regulator in time linear in the time horizon T and cubic in the dimensionality of the state space D. The constrained problem is solved using dual gradient descent, which uses LQG as a subroutine to solve for the primal variables in each iteration and increments the dual variable on the constraint until it is satisfied.

Updating θ is straightforward, since expectations taken with respect to the trajectory induced by π are always conditioned on s_t and all outer expectations over s_t are taken with respect to the trajectory induced by ψ. Therefore, π is essentially decoupled from the transition probability $p\left(s_{t+1} \middle| s_t, a_t, t\right)$ and so its parameters can be updated without affecting the distribution of s_t's. The subproblem that needs to be solved to update θ therefore amounts to a standard supervised learning problem.

Since $\psi\left(a_t \middle| s_t, t; \eta\right)$ and $\pi\left(a_t \middle| s_t; \theta\right)$ are Gaussian, $D\left(\theta, \eta\right)$ can be computed analytically. More concretely, if we assume Σ^π to be fixed for simplicity, the subproblem that is solved for updating $\theta = (\omega, \Sigma^\pi)$ is:

$$\min_\theta \mathbb{E}_\psi \left[\sum_{t=0}^{T} \lambda_t^T \mu_\omega^\pi(o_t) + \frac{\nu_t}{2} \left(\operatorname{tr}\left(G_t^{-1} \Sigma^\pi\right) - \log |\Sigma^\pi| \right) \right.$$
$$\left. + \frac{\nu_t}{2} \left(\mu_\omega^\pi(o_t) - \mathbb{E}_\psi\left[a_t \middle| s_t, t\right]\right)^T G_t^{-1} \left(\mu_\omega^\pi(o_t) - \mathbb{E}_\psi\left[a_t \middle| s_t, t\right]\right) \right]$$

Note that the last term is the squared Mahalanobis distance between the mean actions of ψ and π at time step t, which is intuitive as we would like to encourage π to match ψ.

3.4. Convolutional GPS

The problem of learning high-dimensional optimization algorithms presents challenges for reinforcement learning algorithms due to high dimensionality of the state and action spaces. For example, in the case of GPS, because the running time of LQG is cubic in dimensionality of the state space, performing policy search even in the simple class of linear-Gaussian policies would be prohibitively expensive when the dimensionality of the optimization problem is high.

Fortunately, many high-dimensional optimization problems have underlying structure that can be exploited. For example, the parameters of neural nets are equivalent up to permutation among certain coordinates. More concretely, for fully connected neural nets, the dimensions of a hidden layer and the corresponding weights can be permuted arbitrarily without changing the function they compute. Because permuting the dimensions of two adjacent layers can permute the weight matrix arbitrarily, an optimization algorithm should be invariant to permutations of the rows and columns of a weight matrix. A reasonable prior to impose is that the algorithm should behave in the same manner on all coordinates that correspond to entries in the same matrix. That is, if the values of two coordinates in all current and past gradients and iterates are identical, then the step vector produced by the algorithm should have identical values in these two coordinates. We will refer to the set of coordinates on which permutation invariance is enforced as a coordinate group. For the purposes of learning an optimization algorithm for neural nets, a natural choice would be to make each coordinate group correspond to a weight matrix or a bias vector. Hence, the total number of coordinate groups is twice the number of layers, which is usually fairly small.

In the case of GPS, we impose this prior on both ψ and π. For the purposes of updating η, we first impose a block-diagonal structure on the parameters A_t, B_t and F_t of the fitted transition probability density $\tilde{p}\left(s_{t+1} \middle| s_t, a_t, t; \zeta\right) = \mathcal{N}(A_t s_t + B_t a_t + c_t, F_t)$, so that for each coordinate in the optimization problem, the dimensions of s_{t+1} that correspond to the coordinate only depend on the dimensions of s_t and a_t that correspond to the same coordinate. As a result, $\tilde{p}\left(s_{t+1} \middle| s_t, a_t, t; \zeta\right)$ decomposes into multiple independent probability densities $\tilde{p}^j\left(s_{t+1}^j \middle| s_t^j, a_t^j, t; \zeta^j\right)$, one for each coordinate j. Similarly, we also impose a block-diagonal structure on C_t for fitting $\tilde{c}(s_t)$ and on the parameter matrix of the fitted model for $\pi\left(a_t \middle| s_t; \theta\right)$. Under these assumptions, K_t and G_t are guaranteed to be block-diagonal as well. Hence, the Bregman divergence penalty term, $D\left(\eta, \theta\right)$ decomposes into a sum of Bregman divergence terms, one for each coordinate.

We then further constrain dual variables λ_t, sub-vectors of parameter vectors and sub-matrices of parameter matrices corresponding to each coordinate group to be identical across the group. Additionally, we replace the weight ν_t on $D\left(\eta, \theta\right)$ with an individual weight on each Bregman divergence term for each coordinate group. The problem then decomposes into multiple independent subproblems, one for each coordinate group. Because the dimensionality of the state subspace corresponding to each coordinate is constant, LQG can be executed on each subproblem much more efficiently.

Similarly, for π, we choose a $\mu_\omega^\pi(\cdot)$ that shares parameters across different coordinates in the same group. We also impose a block-diagonal structure on Σ^π and constrain the appropriate sub-matrices to share their entries.

3.5. Features and Policy Model

We describe the features $\Phi(\cdot)$ and $\Psi(\cdot)$ at time step t, which define the state s_t and observation o_t respectively.

Because of the stochasticity of gradients and objective values, the state features $\Phi(\cdot)$ are defined in terms of summary statistics of the history of iterates $\left\{x^{(i)}\right\}_{i=0}^{t}$, gradi-

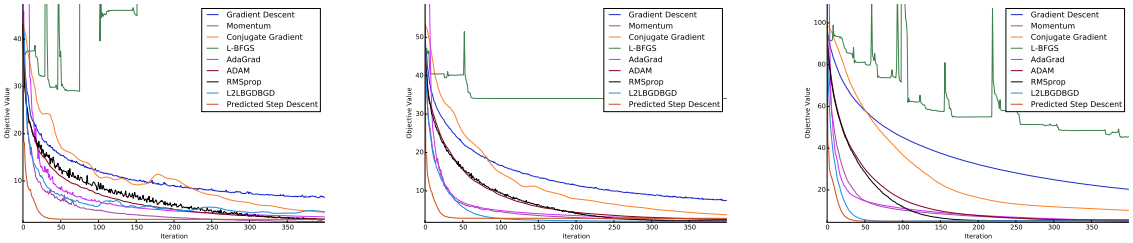

(a) (b) (c)

Figure 1. Comparison of the various hand-engineered and learned algorithms on training neural nets with 48 input and hidden units on (a) TFD, (b) CIFAR-10 and (c) CIFAR-100 with mini-batches of size 64. The vertical axis is the true objective value and the horizontal axis represents the iteration. Best viewed in colour.

ents $\left\{\nabla \hat{f}(x^{(i)})\right\}_{i=0}^{t}$ and objective values $\left\{\hat{f}(x^{(i)})\right\}_{i=0}^{t}$. We define the following statistics, which we will refer to as the average recent iterate, gradient and objective value respectively:

- $\overline{x^{(i)}} := \frac{1}{\min(i+1,3)} \sum_{j=\max(i-2,0)}^{i} x^{(j)}$
- $\overline{\nabla \hat{f}(x^{(i)})} := \frac{1}{\min(i+1,3)} \sum_{j=\max(i-2,0)}^{i} \nabla \hat{f}(x^{(j)})$
- $\overline{\hat{f}(x^{(i)})} := \frac{1}{\min(i+1,3)} \sum_{j=\max(i-2,0)}^{i} \hat{f}(x^{(j)})$

The state features $\Phi(\cdot)$ consist of the relative change in the average recent objective value, the average recent gradient normalized by the magnitude of the a previous average recent gradient and a previous change in average recent iterate relative to the current change in average recent iterate:

- $\left\{\left(\overline{\hat{f}(x^{(t-5i)})} - \overline{\hat{f}(x^{(t-5(i+1))})}\right) / \overline{\hat{f}(x^{(t-5(i+1))})}\right\}_{i=0}^{24}$
- $\left\{\overline{\nabla \hat{f}(x^{(t-5i)})} / \left(\left|\overline{\nabla \hat{f}(x^{(\max(t-5(i+1),t \bmod 5))})}\right| + 1\right)\right\}_{i=0}^{25}$
- $\left\{\frac{\left|\overline{x^{(\max(t-5(i+1),t \bmod 5+5))}} - \overline{x^{(\max(t-5(i+2),t \bmod 5))}}\right|}{\left|\overline{x^{(t-5i)}} - \overline{x^{(t-5(i+1))}}\right| + 0.1}\right\}_{i=0}^{24}$

Note that all operations are applied element-wise. Also, whenever a feature becomes undefined (i.e.: when the time step index becomes negative), it is replaced with the all-zeros vector.

Unlike state features, which are only used when training the optimization algorithm, observation features $\Psi(\cdot)$ are used both during training and at test time. Consequently, we use noisier observation features, which depend directly on the noisy gradients and objective values. The observation features include the following:

- $\left(\hat{f}(x^{(t)}) - \hat{f}(x^{(t-1)})\right) / \hat{f}(x^{(t-1)})$
- $\nabla \hat{f}(x^{(t)}) / \left(\left|\nabla \hat{f}(x^{(\max(t-1,0))})\right| + 1\right)$
- $\frac{\left|x^{(\max(t-1,1))} - x^{(\max(t-2,0))}\right|}{\left|x^{(t)} - x^{(t-1)}\right| + 0.1}$

We use an recurrent neural net with a single layer of 128 LSTM (Hochreiter & Schmidhuber, 1997) cells to model the mean of the policy, $\mu_{\omega}^{\pi}(o_t)$. As mentioned above, the policy operates on each coordinate individually and the parameters of the policy on all coordinates within a coordinate group are identical. Technically, because $\mu_{\omega}^{\pi}(\cdot)$ should only depend on o_t, $\Psi(\cdot)$ also includes the previous state of the LSTM cells.

4. Experiments

For clarity, we will refer to training of the optimization algorithm as "meta-training" to differentiate it from base-level training, which will simply be referred to as "training".

We meta-trained an optimization algorithm on a single objective function, which corresponds to the problem of training a two-layer neural net with 48 input units, 48 hidden units and 10 output units on a randomly projected and normalized version of the MNIST training set with dimensionality 48 and unit variance in each dimension. We used a time horizon of 400 iterations and a mini-batch size of 64 for computing stochastic gradients and objective values. We evaluate the optimization algorithm on its ability to generalize to unseen objective functions, which correspond to the problems of training neural nets on different tasks/datasets. We evaluate the learned optimization algorithm on three datasets, the Toronto Faces Dataset (TFD), CIFAR-10 and CIFAR-100. These datasets are chosen for their very different characteristics from MNIST and each other: TFD contains 3300 grayscale images that have relatively little variation and has seven different categories, whereas CIFAR-100 contains 50,000 colour images that

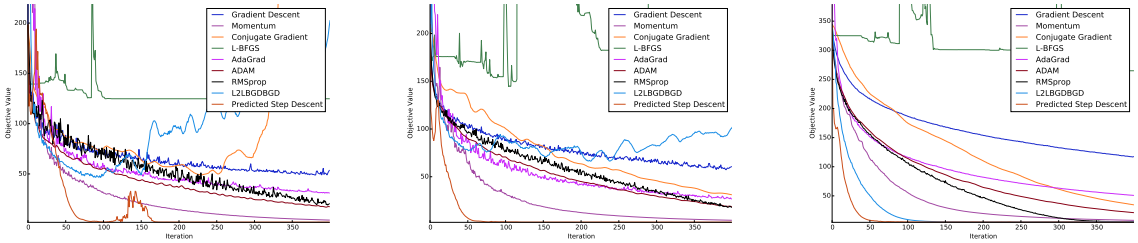

(a) (b) (c)

Figure 2. Comparison of the various hand-engineered and learned algorithms on training neural nets with 100 input units and 200 hidden units on (a) TFD, (b) CIFAR-10 and (c) CIFAR-100 with mini-batches of size 64. The vertical axis is the true objective value and the horizontal axis represents the iteration. Best viewed in colour.

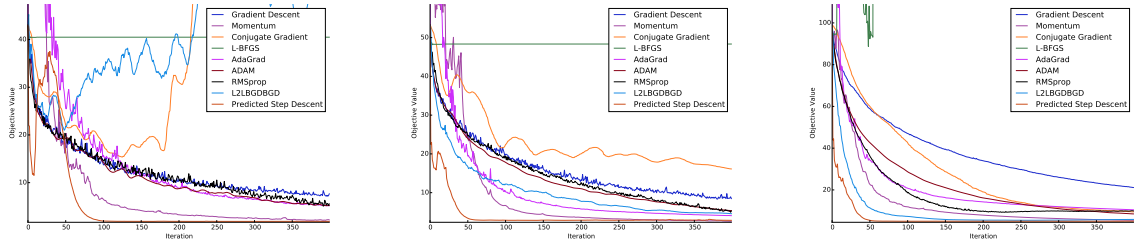

(a) (b) (c)

Figure 3. Comparison of the various hand-engineered and learned algorithms on training neural nets with 48 input and hidden units on (a) TFD, (b) CIFAR-10 and (c) CIFAR-100 with mini-batches of size 10. The vertical axis is the true objective value and the horizontal axis represents the iteration. Best viewed in colour.

have varied appearance and has 100 different categories.

All algorithms are tuned on the training objective function. For hand-engineered algorithms, this entails choosing the best hyperparameters; for learned algorithms, this entails meta-training on the objective function. We compare to the seven hand-engineered algorithms: stochastic gradient descent, momentum, conjugate gradient, L-BFGS, ADAM, AdaGrad and RMSprop. In addition, we compare to an optimization algorithm meta-trained using the method described in (Andrychowicz et al., 2016) on the same training objective function (training two-layer neural net on randomly projected and normalized MNIST) under the same setting (a time horizon of 400 iterations and a mini-batch size of 64).

First, we examine the performance of various optimization algorithms on similar objective functions. The optimization problems under consideration are those for training neural nets that have the same number of input and hidden units (48 and 48) as those used during meta-training. The number of output units varies with the number of categories in each dataset. We use the same mini-batch size as that used during meta-training. As shown in Figure 1, the optimization algorithm meta-trained using our method (which we will refer to as Predicted Step Descent) consistently descends to the optimum the fastest across all datasets. On the other hand, other algorithms are not as consistent and the relative ranking of other algorithms varies by dataset. This suggests that Predicted Step Descent has learned to be robust to variations in the data distributions, despite being trained on only one objective function, which is associated with a very specific data distribution that characterizes MNIST. It is also interesting to note that while the algorithm meta-trained using (Andrychowicz et al., 2016) (which we will refer to as L2LBGDBGD) performs well on CIFAR, it is unable to reach the optimum on TFD.

Next, we change the architecture of the neural nets and see if Predicted Step Descent generalizes to the new architecture. We increase the number of input units to 100 and the number of hidden units to 200, so that the number of parameters is roughly increased by a factor of 8. As shown in Figure 2, Predicted Step Descent consistently outperforms other algorithms on each dataset, despite having not been trained to optimize neural nets of this architecture. Interestingly, while it exhibited a bit of oscillation initially on TFD and CIFAR-10, it quickly recovered and overtook other al-

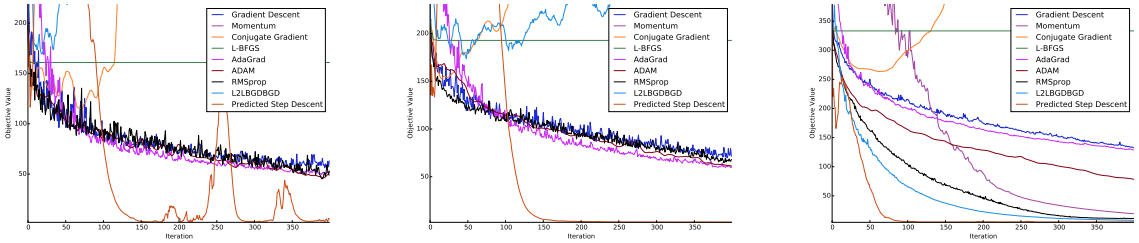

(a) (b) (c)

Figure 4. Comparison of the various hand-engineered and learned algorithms on training neural nets with 100 input units and 200 hidden units on (a) TFD, (b) CIFAR-10 and (c) CIFAR-100 with mini-batches of size 10. The vertical axis is the true objective value and the horizontal axis represents the iteration. Best viewed in colour.

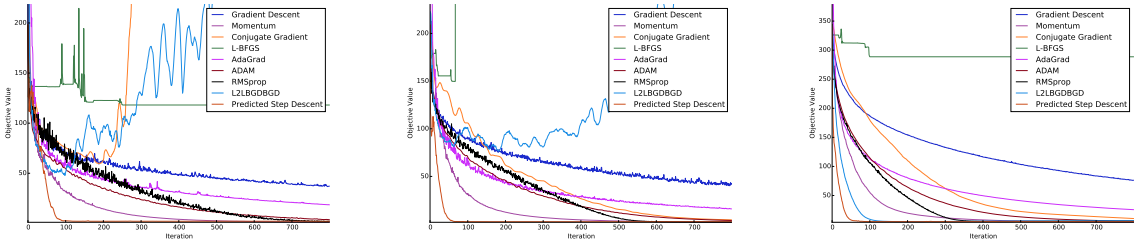

(a) (b) (c)

Figure 5. Comparison of the various hand-engineered and learned algorithms on training neural nets with 100 input units and 200 hidden units on (a) TFD, (b) CIFAR-10 and (c) CIFAR-100 for 800 iterations with mini-batches of size 64. The vertical axis is the true objective value and the horizontal axis represents the iteration. Best viewed in colour.

gorithms, which is reminiscent of the phenomenon reported in (Li & Malik, 2016) for low-dimensional optimization problems. This suggests that it has learned to detect when it is performing poorly and knows how to change tack accordingly. L2LBGDBGD experienced difficulties on TFD and CIFAR-10 as well, but slowly diverged.

We now investigate how robust Predicted Step Descent is to stochasticity of the gradients. To this end, we take a look at its performance when we reduce the mini-batch size from 64 to 10 on both the original architecture with 48 input and hidden units and the enlarged architecture with 100 input units and 200 hidden units. As shown in Figure 3, on the original architecture, Predicted Step Descent still outperforms all other algorithms and is able to handle the increased stochasticity fairly well. In contrast, conjugate gradient and L2LBGDBGD had some difficulty handling the increased stochasticity on TFD and to a lesser extent, on CIFAR-10. In the former case, both diverged; in the latter case, both were progressing slowly towards the optimum.

On the enlarged architecture, Predicted Step Descent experienced some significant oscillations on TFD and CIFAR-10, but still managed to achieve a much better objective value than all the other algorithms. Many hand-engineered algorithms also experienced much greater oscillations than previously, suggesting that the optimization problems are inherently harder. L2LBGDBGD diverged fairly quickly on these two datasets.

Finally, we try doubling the number of iterations. As shown in Figure 5, despite being trained over a time horizon of 400 iterations, Predicted Step Descent behaves reasonably beyond the number of iterations it is trained for.

5. Conclusion

In this paper, we presented a new method for learning optimization algorithms for high-dimensional stochastic problems. We applied the method to learning an optimization algorithm for training shallow neural nets. We showed that the algorithm learned using our method on the problem of training a neural net on MNIST generalizes to the problems of training neural nets on unrelated tasks/datasets like the Toronto Faces Dataset, CIFAR-10 and CIFAR-100. We also demonstrated that the learned optimization algorithm is robust to changes in the stochasticity of gradients and the neural net architecture.

References

Andrychowicz, Marcin, Denil, Misha, Gomez, Sergio, Hoffman, Matthew W, Pfau, David, Schaul, Tom, and de Freitas, Nando. Learning to learn by gradient descent by gradient descent. *arXiv preprint arXiv:1606.04474*, 2016.

Baxter, Jonathan, Caruana, Rich, Mitchell, Tom, Pratt, Lorien Y, Silver, Daniel L, and Thrun, Sebastian. NIPS 1995 workshop on learning to learn: Knowledge consolidation and transfer in inductive systems. https://web.archive.org/web/20000618135816/http://www.cs.cmu.edu/afs/cs.cmu.edu/user/caruana/pub/transfer.html, 1995. Accessed: 2015-12-05.

Bengio, Yoshua. Gradient-based optimization of hyperparameters. *Neural computation*, 12(8):1889–1900, 2000.

Bergstra, James and Bengio, Yoshua. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13(1):281–305, 2012.

Bergstra, James S, Bardenet, Rémi, Bengio, Yoshua, and Kégl, Balázs. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, pp. 2546–2554, 2011.

Bray, Matthieu, Koller-meier, Esther, Müller, Pascal, Van Gool, Luc, and Schraudolph, Nicol N. 3d hand tracking by rapid stochastic gradient descent using a skinning model. In *1st European Conference on Visual Media Production (CVMP*. Citeseer, 2004.

Brazdil, Pavel, Carrier, Christophe Giraud, Soares, Carlos, and Vilalta, Ricardo. *Metalearning: applications to data mining*. Springer Science & Business Media, 2008.

Brazdil, Pavel B, Soares, Carlos, and Da Costa, Joaquim Pinto. Ranking learning algorithms: Using ibl and meta-learning on accuracy and time results. *Machine Learning*, 50(3):251–277, 2003.

Cramer, Nichael Lynn. A representation for the adaptive generation of simple sequential programs. In *Proceedings of the First International Conference on Genetic Algorithms*, pp. 183–187, 1985.

Daniel, Christian, Taylor, Jonathan, and Nowozin, Sebastian. Learning step size controllers for robust neural network training. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

Domke, Justin. Generic methods for optimization-based modeling. In *AISTATS*, volume 22, pp. 318–326, 2012.

Duchi, John, Hazan, Elad, and Singer, Yoram. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.

Feurer, Matthias, Springenberg, Jost Tobias, and Hutter, Frank. Initializing bayesian hyperparameter optimization via meta-learning. In *AAAI*, pp. 1128–1135, 2015.

Fu, Jie, Lin, Zichuan, Liu, Miao, Leonard, Nicholas, Feng, Jiashi, and Chua, Tat-Seng. Deep q-networks for accelerating the training of deep neural networks. *arXiv preprint arXiv:1606.01467*, 2016.

Graves, Alex, Wayne, Greg, and Danihelka, Ivo. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.

Gregor, Karol and LeCun, Yann. Learning fast approximations of sparse coding. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 399–406, 2010.

Hansen, Samantha. Using deep q-learning to control optimization hyperparameters. *arXiv preprint arXiv:1602.04062*, 2016.

Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

Hochreiter, Sepp, Younger, A Steven, and Conwell, Peter R. Learning to learn using gradient descent. In *International Conference on Artificial Neural Networks*, pp. 87–94. Springer, 2001.

Hutter, Frank, Hoos, Holger H, and Leyton-Brown, Kevin. Sequential model-based optimization for general algorithm configuration. In *Learning and Intelligent Optimization*, pp. 507–523. Springer, 2011.

Kingma, Diederik and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Levine, Sergey, Finn, Chelsea, Darrell, Trevor, and Abbeel, Pieter. End-to-end training of deep visuomotor policies. *arXiv preprint arXiv:1504.00702*, 2015.

Li, Ke and Malik, Jitendra. Learning to optimize. *CoRR*, abs/1606.01885, 2016.

Liang, Percy, Jordan, Michael I, and Klein, Dan. Learning programs: A hierarchical Bayesian approach. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 639–646, 2010.

Maclaurin, Dougal, Duvenaud, David, and Adams, Ryan P. Gradient-based hyperparameter optimization through reversible learning. *arXiv preprint arXiv:1502.03492*, 2015.

Ruvolo, Paul L, Fasel, Ian, and Movellan, Javier R. Optimization on a budget: A reinforcement learning approach. In *Advances in Neural Information Processing Systems*, pp. 1385–1392, 2009.

Schmidhuber, Jürgen. Optimal ordered problem solver. *Machine Learning*, 54(3):211–254, 2004.

Snoek, Jasper, Larochelle, Hugo, and Adams, Ryan P. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pp. 2951–2959, 2012.

Sprechmann, Pablo, Litman, Roee, Yakar, Tal Ben, Bronstein, Alexander M, and Sapiro, Guillermo. Supervised sparse analysis and synthesis operators. In *Advances in Neural Information Processing Systems*, pp. 908–916, 2013.

Swersky, Kevin, Snoek, Jasper, and Adams, Ryan P. Multi-task bayesian optimization. In *Advances in neural information processing systems*, pp. 2004–2012, 2013.

Thrun, Sebastian and Pratt, Lorien. *Learning to learn*. Springer Science & Business Media, 2012.

Tieleman, Tijmen and Hinton, Geoffrey. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2), 2012.

Vilalta, Ricardo and Drissi, Youssef. A perspective view and survey of meta-learning. *Artificial Intelligence Review*, 18(2):77–95, 2002.

Wang, Huahua and Banerjee, Arindam. Bregman alternating direction method of multipliers. *CoRR*, abs/1306.3203, 2014.